Rapid Updates of GIS Databases from Digital Images

Howard Schultz, Allen R. Hanson, Edward M. Riseman, Frank Stolle
University of Massachusetts
Computer Science Department
Amherst, MA 01003
{schultz|hanson|riseman|stolle@cs.umass.edu}
http://vis-www.cs.umass.edu

Abstract

Our goal is to improve the way in which Geographic Information System (GIS) databases are created, updated and utilized. We are building software systems that will enable users to rapidly and inexpensively acquire, update, analyze and visualize high-resolution digital elevation maps (DEM), ortho-images, and other GIS products from digital images collected by a wide range of aerial and satellite platforms. The low cost and short turnaround time for these products will enable education and research institutes, government agencies and commercial organizations to expand the scope of their GIS applications, especially in the areas of environmental monitoring, change detection, and urban planning.

1 Introduction

Rapid improvements in the spatial, temporal and spectral resolution of image sensors have significantly increased the quality and quantity of image data while decreasing the acquisition cost. These advances, coupled with recent progress in computer vision techniques for reducing and organizing geospatial data, have the potential to revolutionize the way in which Geographic Information System (GIS) are created, updated and utilized. However, these technical advances have come at the cost of an explosive increase in the bandwidth required to transmit the data and the storage capacity needed to retain the data. To reduce the demands on transmission and storage systems, as well as reduce the time and expense required to process the data, our program is focused on developing highly efficient and automatic methods for generating geographically registered terrain models.

Our approach reduces a highly redundant, overlapping sequence of images to a non-overlapping sequence of terrain model tiles. The system employs an integrated approach to minimizing sensor pose, geographic registration, and DEM generation errors. The system incorporates a unique solution to sensor pose recovery that combines position and orientation data (obtained from GPS and INS instruments) (Schultz, et al., 1999a) and relative orientation information (derived from the image sequence) in a Kalman-Bucy filter. The improved pose estimate enables us to extract accurate shape and geospatial uncertainty information from an image sequence. The system incorporates the DEM extraction tool, Terrest (Schultz, 1994 and 1995), which has several characteristics that make it well suited for motion imagery applications. These features include stability with respect to oblique viewing angles, optimized numerical code, and the ability to estimate the geospatial uncertainty of the recovered DEM. The inclusion of geospatial uncertainty in terrain models gives the system a wide range of applications, including change detection and DEM recovery over steep terrain.

The principal challenge is to define and implement a unified approach to terrain modeling that encompasses the entire process, including designing the instrumentation package, sensor pose recovery, calibration, softcopy photogrammetry, and error propagation.

2 A Photogrammetric Approach.

As an object moves through the field-of-view of an aerial camera, it is seen from a continuously changing viewpoint. The motion parallax captured by the sequence of images contains information about the 3D structure of the object. The 3D structure of the surface can be seen by constructing a stereoscopic pair of mosaics, one assembled from the leading edge, and one from the trailing edge of each frame (Zhu, et at. 1999a, b). Because mosaics can be produced in real-time, the terrain's topography can be viewed immediately using polarized or anaglyph viewing systems. The stereoscopic pair of mosaics and a red-blue overlay is shown are Figure 1.

When the mosaic approach is extended to form accurate geographically registered 3D terrain models, certain difficulties emerge: (i) Deviation from a straight flight path (caused by cross winds, turbulence, control inaccuracies, etc.) results in a form of geometric distortion characterized by curved epipolar lines, which significantly increases the processing time required to generate a disparity map from a stereoscopic mosaic pair. (ii) The abrupt change as individual image strips are added to the mosaic creates seams, which may interfere with the image matching process. (iii) Registering a mosaic to an existing DEM requires ground control points (GCPs), which are difficult to acquire and are often inaccurate. Furthermore, the terrain model must be interpolated between GCPs.

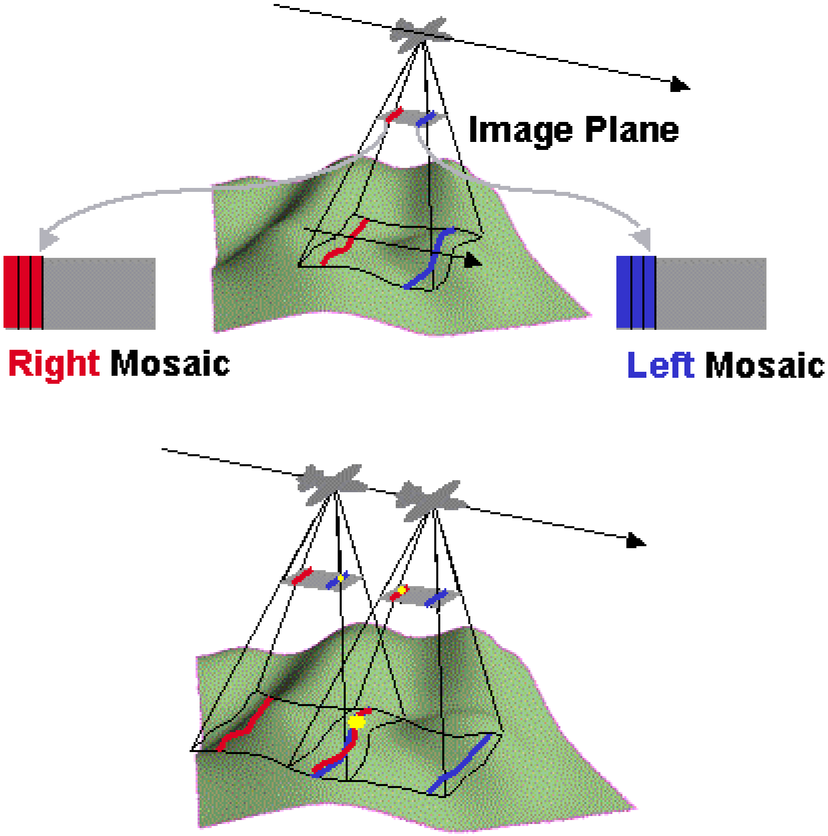

Figure 1. (Top) Two mosaics are generated for a single down-looking camera by stitching together slices from the leading and trailing edges of each frame. (Bottom) Points on the surface are imaged twice, once in the leading edge mosaic and once in the trailing edge mosaic.

For applications requiring more precise geo-registered terrain models, we propose using a direct method based on classic photogrammetry and optimal pose estimation. This approach produces a series of small geo-registered 3D terrain model tiles, where each tile consists of a DEM, ortho-image, and geospatial uncertainty map. The terrain model tiles are then pieced together to form a larger composite terrain model. Because the individual terrain models are generated directly from image pairs, this approach eliminates the curved epipolar line problem. Furthermore, model warping is not required because the tiles are in a world coordinate system.

In photogrammetric terms, the tiles are in a 3D *model space*, which is related to the terrain 3D *object space* by a rigid body rotation, translation and scale change. Furthermore, determination of the parameters of the model space (recovering the epipolar geometry) requires only image-to-image tie-points (which are generated during the real-time mosaic generation process). In other words, a 3D model is created without any GCPs. Of course, registering the 3D model to the terrain does require GCPs.

This approach requires an ability to consistently generate terrain models in real-time. Here real-time refers to an ability to produce terrain model tiles at a rate that keeps up with the aircraft flying over the terrain. We are not referring to video frame rates. The frame rate will vary depending on the aircraft speed, altitude, image format and desired ground sampling distance. For a lightweight aircraft flying at 50 m/sec, typical frame rates vary from once per second to once every eight seconds. Because DEM generation and geo-registration depend on accurate orientation and location information at the time of exposure, an important objective is to generate accurate camera pose estimates in real-time from standard navigation instruments.

3 Real-time pose estimation

The use of standard Kalman-Bucy filtering (KBF) techniques can provide estimates of an aircraft's location and orientation. A Kalman filter forms the heart of the standard approach of navigation systems. Such systems typically use Inertial Navigation Systems (INS) to provide information regarding the aircraft orientation and position. The INS measurement system is often augmented by other external navigation aids (such as Long Range Navigation (LORAN), Terrain Contour Matching (TERCOM), or the Global Positioning System (GPS)) that can provide additional information about the aircraft.

Photogrammetry provides an additional measurement of aircraft motion. The image-to-image tie-points found during mosaic generation are used in conjunction with a bundle adjustment (BA) procedure to compute the relative, exterior orientation, which gives the change in position and orientation from one frame to the next. The output of the BA can then be used to estimate the rate of change of the aircraft position and orientation by simply dividing the relative, exterior orientation parameters by the time between frames. These results are combined with GPS and INS measurements and a simple KBF motion model to improve estimates of the aircraft position and orientation. Our preliminary results show that the addition of relative, exterior orientation information to the usual navigation measurements reduce position and orientation errors by at least 33%.

4 DEM estimation

We use the UMass Terrain Reconstruction System[1] (Terrest) to recover topography from a sequence of overlapping digital images (Schultz, 1994, 1995). The highly redundant information captured by a sequence of images enables Terrest to estimate an optimal DEM, and geospatial uncertainty map (error bars of the elevation estimates). Terrest is organized in two parts - a collection of softcopy

---

[1] Terrest was selected by NIMA as part of the Pathfinder 2000 AFE/ATR program (http://www.discover-aai.com/pathfinder.htm).

photogrammetric routines (rectify, matching and ortho-rectify), which generates a DEM and ortho-image from a pair of images, and DEM fusion, which computes an optimal DEM and geospatial uncertainty map from a sequence of DEMs.

The DEM fusion procedure uses the concept of self-consistency (Leclerc, et al., 1998a,b) to identify errors in the DEM. We have shown that it is highly unlikely that elevation mistakes caused by matching errors (which are sometimes referred to as blunders) will be found in the same physical location. As a result, a series of overlapping DEMs tends to agree for correct matches and disagree for false matches. This property enables the algorithm to identify blunders by looking for large discrepancies between the DEMs (Schultz, et al. 1999b).

Consider the example of the four overlapping images shown in Figure 2. Starting with four images, twelve DEMs were computed. Based on the self-consistency principle, a threshold was selected that rejected any elevation estimates that deviated from the mean of the distribution by two or more standard deviations. Next, the optimal elevations and the geospatial errors were found by taking the average and standard deviation of the elevation estimates that passed the threshold test. A rendered view of the optimal DEM, as well as an image showing the number of contributing estimates, are shown in Figure 3. Although no individual DEM was free from errors, 99.5% of the surface points in the optimal DEM had at least 6 estimates that passed the self-consistency test. For this example, the ground sampling distance was approximately 35cm, and the RMS elevation error was 17cm over an elevation range of 123m.

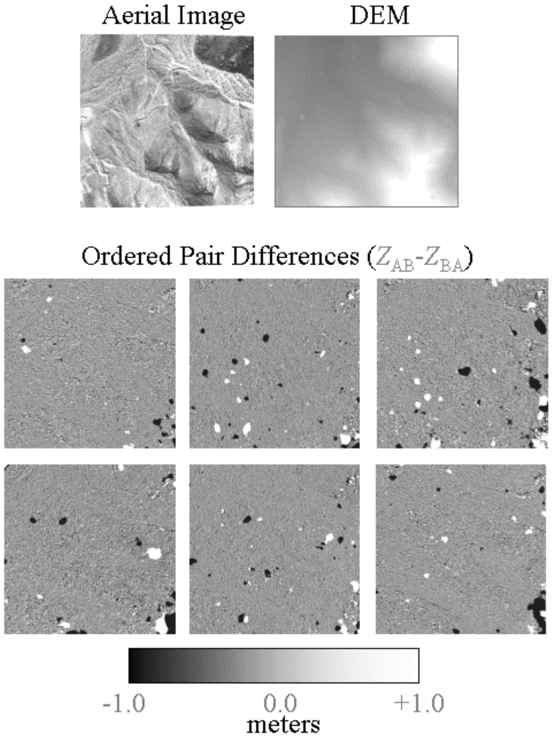

Figure 2. (Top) One of the four aerial images and one of the 12 DEMs used in the experiment. (Bottom) The six ordered pair differences created from the twelve DEMs. Black and white regions indicate areas of inconsistent elevations.

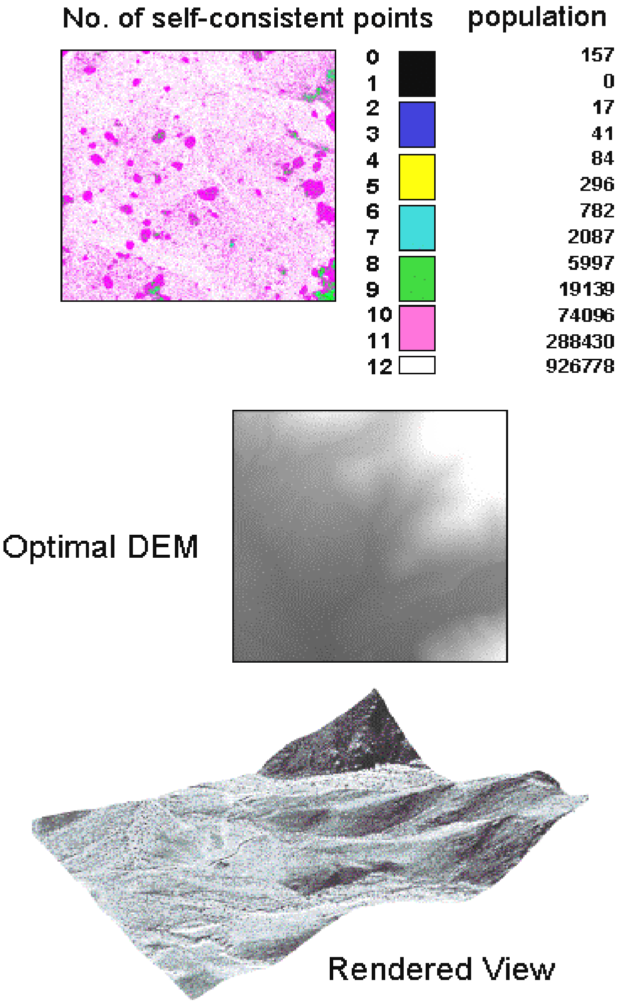

Figure 3. The results of self-consistency analysis. A map of the number of self-consistent points (top); the optimal DEM made by averaging all self-consistent elevation estimates (center); and a 3D rendering of the optimal DEM (bottom). 99.5% of the grid points in the DEM had at least 6 good (self-consistent) elevation estimates. The ground sampling distance was approximately 35cm, and the RMS elevation error was 17cm over an elevation range of 123m.

5 Towards a real-time system

A diagram of a proposed real-time terrain modeling system is shown in Figure 4. Each block represents an independent process, which is invoked when a sufficient amount of input data become available. The real-time pose recovery algorithm assimilates output from the GPS, INS and the bundle adjustment. The block matching process provides data (in the form of image-to-image tie-points) to the stereoscopic mosaic generation and the bundle adjustment modules. The Kalman-Bucy filtering algorithm has two parts, a predictor which estimates the position and orientation of the image sensor at the next exposure, and a corrector which computes the optimal values of the position and orientation of the sensor a few seconds in the past. The predicted sensor pose initializes the block matching algorithm, which will improve its speed and accuracy. The corrected pose is used by Terrest to compute the terrain model tiles. The tiles are then assembled to form a larger terrain model.

By far, the Terrest modules are the most time consuming procedures. Tests using a DEC Alpha server with four 600MHz processors show that Terrest runs about 50 times slower than real-time. Thus, the key to a successful system rests on the availability of high-speed onboard computing systems. Shifting some of the data reduction functions from the ground to the aerial platform can reduce the demands on the communication links, allow for more efficient use of the captured data, and improve the real-time estimates of the aircraft's navigation information.

6 Summary

The techniques described in this paper have been successfully applied to a wide variety of academic research, commercial and defense projects. We believe that the integration of these techniques into an operational system will substantially increase the value of image databases and improve the capabilities of many GIS applications. In parallel efforts, we are developing a real-time stereoscopic mosaic visualization, and a 3D terrain modeling system based on the photogrammetric approach discussed above.

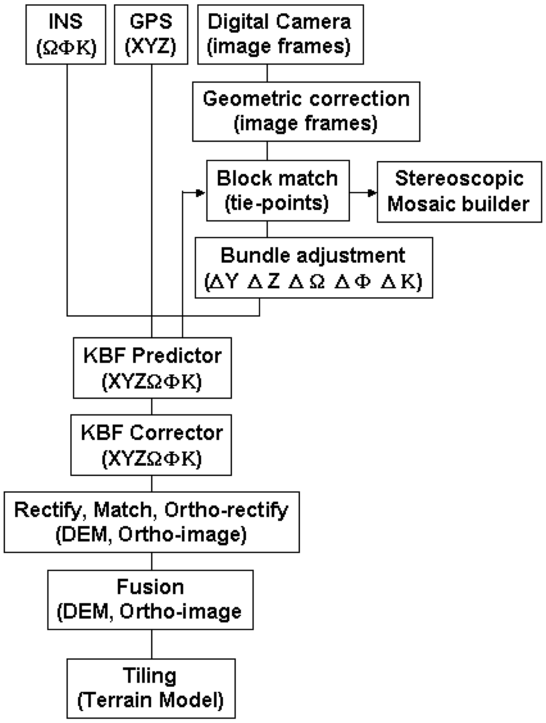

Figure 4. A block diagram of a real-time system. Each box indicates an independent process, with the output data products shown in parentheses. The data flow is from top to bottom, except where indicated with an arrow. The top row shows the data sources. After correcting for lens distortion, the block match algorithm finds tie-points between the latest image frame and the previous one. The tie-points are used by the stereo mosaic builder and bundle adjustment module, which computes the rate of change of the camera pose. The pose and change in pose information are fed to the KBF module, which consists of two parts - a predictor and a corrector. The predicted pose is used to constrain the block matching process. The corrected pose is used by the terrain modeling modules (rectify, match, ortho-rectify and fuse) to compute terrain model tiles. The individual tiles are then assembled to form a larger terrain model.

Acknowledgements

This research has been supported by the Army Research Office under grant number DAAD19-99-1-0016, the National Science Foundation under grants EIA-9726401, EIA-9726401 (KOSEF) and EIA-0105272 (SGER).

References

Leclerc, Y. G., Q. T. Luong and P. Fua, “A Framework for Detecting Changes in Terrain,” *IEEE Trans. On Pattern Analyssi and Machine Intel.* Vol. 20., no. 11., pp 1143-1160, November 1998.

Leclerc, Y. G., Q. T. Luong and P. Fua, “Self-consistency: A Novel Approach to Characterizing the Accuracy and Reliability of Point, Correspondences Algorithms,” DARPA Image Understanding Workshop, Monterey CA, pub by Morgan Kauffman, 1998

Schultz, H., Slaymaker, D., Hanson, A., Riseman, E., Holmes, C., Powell, M., and Delaney, M., "Cost-effective Determination of Biomass from Aerial Images," *International Workshop on Integrated Spatial Data*, Portland, Maine, June 14-15, 1999.

Schultz, H., Terrain Reconstruction from Oblique Views, Proc. *DARPA Image Understanding Workshop*, Monterey, CA, pp. 1001-1008, 1994.

Schultz, H., “Terrain Reconstruction from Widely Separated Images”, Proc. SPIE, Volume 2486, pp. 113-123, Orlando, FL, April, 1995.

Schultz, H., Schwickerath, A., Stolle, F., Hanson, A., Riseman, E., "Incremental Digital Elevation Map Generation From stereo Images," ISPRS Workshop on Theoretical and Practical Aspects of Surface Reconstruction and 3D Object Extraction, Haifa, Israel, September 9-11,1997.

Schultz, H., Hanson, A., Holmes, C., Riseman, E., Slaymaker, D., Stolle, F., “Integrating Small Format Aerial Photography, Videography, and a Laser Profiler for Environmental Monitoring”, *ISPRS WG III/1 Workshop on "Integrated Sensor Calibration and Orientation*, Portland, Maine, USA June 16-17, 1999.

Schultz, H., Riseman, E., Stolle, F. Error Detection and DEM Fusion Using Self-Consistency, Submitted to the *7th IEEE International Conference on Computer Vision*, Kerkyra, Greece. 1999.

Zhu, Z., Hanson, A R., Schultz, H., Stolle, F., Riseman, E. M., Stereo Mosaics from a Moving Video Camera for Environmental Monitoring, First International Workshop on Digital and Computational Video, pp. 45-54, Tampa, Florida, Dec 10, 1999.

Zhu, Z, Riseman, E. M., Hanson, A. R., Schultz, H., Automatic Geo-Correction of Video Mosaics for Environmental Monitoring, Report TR #99-28, Computer Science Department, University of Massachusetts at Amherst, April, 1999